

Illumination in Diverse Codimensions

David C. Banks

Institute for Computer Applications in Science and Engineering
MailStop 132C, NASA Langley Research Center
Hampton, Virginia 23681

ABSTRACT

This paper derives a model of diffuse and specular illumination in arbitrarily large dimensions, based on a few characteristics of material and light in 3-space. It then describes how to adjust for the anomaly of excess brightness in large codimensions. If a surface is grooved or furry, it can be illuminated with a hybrid model that incorporates both the 1D geometry (the grooves or fur) and the 2D geometry (the surface).

Research was supported by the National Aeronautics and Space Administration under NASA Contract No. NAS1-19480 while the author was in residence at the Institute for Computer Applications in Science and Engineering (ICASE), NASA Langley Research Center, Hampton, VA 23681.

1 Introduction

When a geometric object possesses a distinct (outward) unit normal at each point, the familiar models of illumination can be applied to the object. When the object is in a large dimensional space, the space of unit normals has two or more dimensions (that is, the codimension of the object is at least two) and the illumination model must be extended; examples include curves in 3-space and 4-space, or surfaces in 4-space and higher.

This paper addresses the problem of applying light in large codimensions. Consider an object of dimension $k > 0$ in Euclidean space of dimension $n > k$. The difference $n - k$ is the *codimension* of the object. Mathematicians use the term *k-manifold* to denote the k -dimensional generalization of curves (1-manifolds) and surfaces (2-manifolds). Every neighborhood of a k -manifold is homeomorphic with Euclidean k -space.

Regarding codimension 1, popular texts on computer graphics [Foley90] [Rogers85] handle the special case of $k=2$ and $n=3$: these are ordinary surfaces in 3-space. Other authors [Carey87] [Steiner87] have noted that whenever the codimension is 1, each point of a manifold can be naturally assigned a normal vector. The usual lighting equations then prevail. (Special care is required for non-orientable manifolds or manifolds with boundary, since their “frontfacing” elements are not well defined.)

Regarding codimension 2, several authors have considered the case of $k=1$, $n=3$ for illuminating fur [Kajiya85], hair [Anjyo92] [Miller88] [LeBlanc91] [Watanabe92], or anisotropic grooves on a surface [Kajiya89] [Poulin90] [Westin92] [Ward92]. The case $k=2$, $n=4$ has been studied in the “Fourphront” system [Banks92] [Banks93] and also by Hanson [Hanson93] for examining a variety of surfaces in 4-space. Kajiya and Hanson testify that this model is not based on physical principles, calling it an “ad hoc” and “heuristic” result, respectively. But in fact the Kajiya-Hanson model can be derived from a few physical principles. This is the subject of section 2.

In daily life one encounters illuminated surfaces everywhere. It is reasonable to believe that the human visual system is especially well designed to infer shape from the shading of 2-dimensional surfaces in 3-space [Horn89]. If surfaces in 3-space represent the ideal for visual comprehension, the Kajiya-Hanson model suffers from peculiar drawback: manifolds are “too bright” when the codimension grows larger. Section 3 explains the cause and presents a simple remedy to the problem.

Kajiya noted the importance of global illumination effects (in the form of attenuation and shadows) for rendering textured volume elements. Section 4 shows how the combination of a manifold together with a vector field (like a surface together with fur) can be illuminated to simu-

late global effects. The technique can be incorporated into a simple object-order (e.g., polygon) renderer.

2 The Model for Large Codimensions

The final results of this section will be equations for diffuse and specular illumination that are equivalent to the results that Kajiya and Hanson have presented [Kajiya89] [Hanson93]. The new contribution that this section offers is a physical motivation to the derivation. The conventional motivation begins by promoting the dimension of a manifold, illuminating the promoted manifold, and integrating. The new motivation dispenses with the promotion and integration steps altogether. It proceeds directly from the geometry to the illumination solution, without regard to the participating dimensions.

The following discussion makes heavy use of the tangent space \mathbf{T} and the normal space \mathbf{N} at a point \mathbf{p} on a k -manifold M in n -space (see Figure 1). The space \mathbf{T} is the vector space tangent to a point in M . It has dimension k , matching that of the manifold M . The space \mathbf{N} is orthogonal to \mathbf{T} and has dimension c (the codimension of M). The dimensions of \mathbf{T} and \mathbf{N} add up to the dimension n of the entire space, of course.

2.1 Conventional Motivation

The benefit of codimension 1 is that there exist only two unit normals in a point's 1-dimensional normal space. The usual illumination equations require the modest choice of one of the two. If the codimension is large, there is no clear way to select one unit normal from the infinitude that are available. There is a clever solution that other authors have adopted: the dimension of the manifold can be promoted to reduce the codimension.

Let $S^n(r)$ denote an n -sphere of radius r . A circle of radius 10 is then $S^1(10)$; a unit sphere is $S^2(1)$, or simply S^2 . Kajiya, Hanson, and others have proposed that illuminating a k -manifold M of codimension $c > 1$ can be accomplished after forming the cross-product of M with $S^{c-1}(r)$. It is required that $S^{c-1}(r)$ lie within the normal space \mathbf{N} . A point is thus promoted to a circle in 2-space or to a sphere in 3-space; a curve is promoted to a tube in 3-space; a surface is promoted to a volume in 4-space.

The advantage of promoting M to $M' = M \times S^{c-1}(r)$ is that the promoted manifold has codimension 1. This represents the simple case where the usual lighting equations prevail. The promoted manifold M' can provide an effective representation of M with no further processing. But to render M itself, one must employ a scheme whereby a point \mathbf{p} in M inherits the illumination of its fiber $\mathbf{p} \times S^{c-1}(r)$ in M' . A reasonable way to accomplish that goal is to integrate the intensity of the reflected light over $\mathbf{p} \times S^{c-1}(r)$ and then to average it. The average intensity is obtained by dividing

the integrated intensity by the measure of the fiber as seen by the eye. This measure can be a length, an area, a volume, or so forth, in accordance with the dimension $c-1$ of the sphere $S^{c-1}(r)$ (used in the cross product) over which the average is taken. The limit of the average, as $r \rightarrow 0$, yields a reasonable intensity for the point \mathbf{p} .

There are two drawbacks to this approach of promoting M to M' , integrating, and then averaging. First, the integration is unwieldy for $c > 1$, due to the specular term in the integrand. Second, the projected measure of $S^{c-1}(r)$ is view-dependent. This opposes the notion that diffuse reflection is view-independent. For example, in derivation (13) of [Kajiya89], the integrated intensity over a fiber $\mathbf{p} \times S^1(r)$ of M' is calculated to be

$$\begin{aligned} I'_{diffuse} &= k_d r \mathbf{L} \cdot \mathbf{L}_N \int_0^\pi \sin \theta \, d\theta \\ &= k_d 2r \mathbf{L} \cdot \mathbf{L}_N \end{aligned}$$

where k_d is the diffuse coefficient, \mathbf{L} is the light vector, and \mathbf{L}_N is the projection of \mathbf{L} onto \mathbf{N} (Figure 1). Under a parallel projection, the arclength of the circle can vary from $2r$ (viewing the tube from the side) to πr (viewing the tube end-on). So the average intensity ranges between a minimum of $2/\pi k_d \mathbf{L} \cdot \mathbf{L}_N$ and a maximum of $k_d \mathbf{L} \cdot \mathbf{L}_N$ according to the viewing angle. Kajiya avoided this problem by treating the quantity $k_d 2r / \text{projectedArclength}(r)$ as a constant, giving a diffuse quantity of

$$I_{diffuse} = K_d \mathbf{L} \cdot \mathbf{L}_N$$

for a point on the original manifold M .

2.2 Principles for Diffuse Reflection

One can, in fact, justify Kajiya's result by characterizing diffuse reflection in the following way. A neighborhood of a point \mathbf{p} absorbs energy from the incoming light (which delivers I_{source} per unit cross section), and then it re-radiates a fraction k_d of the absorbed energy. How much energy




Figure 1. Light shines in direction \mathbf{L} at a point \mathbf{p} on a tube. \mathbf{L}_N is the projection of the light onto the normal space \mathbf{N} . The diffuse reflection is integrated over the visible portion of the circle S^1 .

does the beam deliver to a unit-neighborhood of \mathbf{p} ? That depends on the cross section of the beam and the angle it makes with the tangent plane (Figure 2).

Suppose an incident light beam strikes M at \mathbf{p} . The light vector \mathbf{L} (pointing in the direction that the beam propagates) projects orthogonally onto the tangent space \mathbf{T} at \mathbf{p} to produce the vector $\mathbf{L_T}$. The two vectors form an angle $a(\mathbf{L}, \mathbf{L_T})$. Simple trigonometry shows that a unit neighborhood of the tangent space intercepts a beam whose cross-section has measure $\sin(a)$. Note that this quantity is never negative, since a vector can be no more than 90° from the tangent space. The manifold re-radiates k_d of the energy delivered by the beam's cross-section. Thus the diffuse component of reflection at \mathbf{p} is given by

$$(I) \quad I_{diffuse} = k_d I_{source} \sin a(\mathbf{L}, \mathbf{L_T}).$$

This solution is essentially the same as Kajiya's: the sine (measured against \mathbf{T}) and cosine (measured against \mathbf{N}) are equal.

The principles for this result are (1) the re-radiated light's intensity varies with the energy delivered by the incident beam; and (2) the manifold re-radiates isotropically.

Equation (I) is purely local, neglecting any effects of shadowing (even self-shadowing). For a closed surface in 3-space, it is common practice to clamp the diffuse term to zero when the surface normal points away from the light source. This is best regarded as a "global" calculation. A very thin surface *does* re-radiate light both forward and backward, as the local model predicts. Moreover, when the codimension is larger than 1, the unit normals form a connected set. In that case there is no "front" or "back" side of the manifold. Local two-sidedness is an exclusive property of codimension one.

2.3 Principles for Specular Reflection

The characteristic of a specular highlight is that it indicates locations on a manifold where the angle between the reflection vector \mathbf{R} and the view vector \mathbf{V} is zero. One can use an exponential function to condense the reflected intensity into the region where this angle is small (Phong




Figure 2. Light shines in direction \mathbf{L} at a point \mathbf{p} . $\mathbf{L_T}$ is the projection of the light onto the tangent space \mathbf{T} . The beam strikes a unit neighborhood of \mathbf{p} at an angle a .

lighting). The only problem is how to determine an appropriate unit-length reflection vector \mathbf{R} when the codimension exceeds 1: there are infinitely many to choose from.

A simple principle to invoke is that light, in a uniform medium, follows a path of (locally) minimum length. To see how the principle applies to a ray of light reflecting from a tangent space \mathbf{T} , consider a point source \mathbf{q} that shines on the point \mathbf{p} and bounces to reach a point \mathbf{u} (Figure 3). The segments from \mathbf{q} to \mathbf{p} and from \mathbf{p} to \mathbf{u} are straight-line paths, so individually they satisfy the minimal-distance criterion.

The total distance must be a local minimum as well. If the points were all in a plane the problem would be very easy: the angles $a = \angle(\mathbf{q} \mathbf{p} \mathbf{q}_T)$ and $b = \angle(\mathbf{u} \mathbf{p} \mathbf{u}_T)$ must be identical, with the tangent projections \mathbf{q}_T and \mathbf{u}_T lying on the opposite sides of \mathbf{p} . The actual situation is nearly this simple. If $\mathbf{u} - \mathbf{p}$ really is a reflection vector then a path from \mathbf{q} to \mathbf{u} via a nearby point $\mathbf{p} + \mathbf{s}$ in \mathbf{T} must be longer than the path via \mathbf{p} . That immediately forces \mathbf{q}_T , \mathbf{p} , and \mathbf{u}_T to be collinear. To see why, consider choosing \mathbf{s} off of the line $\overline{\mathbf{u}_T \mathbf{q}_T}$. If \mathbf{p} lies on $\overline{\mathbf{u}_T \mathbf{q}_T}$ then perturbing the path over to $\mathbf{p} + \mathbf{s}$ increases the base length of each triangle (by the Cauchy-Schwartz inequality), hence increasing each hypotenuse, and hence increasing the total path-length. So the triangles (Figure 3) lie in the planes $(\mathbf{q} \overline{\mathbf{u}_T \mathbf{q}_T})$ and $(\mathbf{u} \overline{\mathbf{u}_T \mathbf{q}_T})$.

Now consider the situation when $b = a$. What happens when \mathbf{p} is perturbed (in the line $\overline{\mathbf{u}_T \mathbf{q}_T}$) over to some $\mathbf{p} + s(\mathbf{q}_T - \mathbf{u}_T)$? The trigonometry is exactly the same as for the “easy” case of the plane. The total distance $D(s)$ is parametrized by s :

$$D(s) = d(\mathbf{q}, \mathbf{p} + s(\mathbf{q}_T - \mathbf{u}_T)) + d(\mathbf{p} + s(\mathbf{q}_T - \mathbf{u}_T), \mathbf{u})$$

A straightforward application of trigonometry and calculus demonstrates that the total distance is a local minimum. One must simply verify that $\frac{d}{ds} D(s) = 0$ when $s = 0$. As a result, $b = a$.

The unit vectors \mathbf{L} and \mathbf{R} consequently have identical tangent components, so the first requirement on a unit reflection vector \mathbf{R} is that $\mathbf{R}_T = \mathbf{L}_T$. If the codimension is 1, there are two such “reflection” vectors, \mathbf{R}^+ and \mathbf{R}^- ; \mathbf{R}^+ is the continuation of \mathbf{L} transmitted through \mathbf{T} (for opaque manifolds of codimension 1, this solution is ignored). When the codimension is 2, the set of all




Figure 3. A ray of light emanates from \mathbf{q} , strikes the tangent space \mathbf{T} at \mathbf{p} , and reflects to \mathbf{u} . If the total path has minimum length, angles a and b are equal.

reflection vectors forms a cone-shaped family R (Figure 4). The unit reflection vectors from R project to a circle in the normal space \mathbf{N} . In general, the unit reflection vectors project to $S^{c-1}(r)$ in \mathbf{N} when the codimension is c .

What is the angle between the view vector \mathbf{V} and the space R of reflections? It is the angle between \mathbf{V} and the closest vector \mathbf{R} in R . This vector is easy to find. A unit reflection \mathbf{R} can be expressed by its tangent and normal components $\mathbf{R}_T = \mathbf{L}_T$ and \mathbf{R}_N . The unit view vector can be likewise decomposed into \mathbf{V}_T and \mathbf{V}_N .

The components \mathbf{V}_T , \mathbf{V}_N , and \mathbf{R}_T are all fixed, so the distance between \mathbf{R} and \mathbf{V} is minimized when $\|\mathbf{R}_N - \mathbf{V}_N\|$ is minimized. That occurs when \mathbf{R}_N and \mathbf{V}_N are collinear: $\mathbf{R}_N = \lambda \mathbf{V}_N$ for some scalar λ . To see why this is minimal, recall that the vector \mathbf{R}_N is also perpendicular to the point \mathbf{R}_N on the sphere S^{c-1} in the normal space. It is a familiar result from calculus that if the distance from a point \mathbf{p} (off of S^{c-1}) to a point \mathbf{q} (on S^{c-1}) is minimal, the vector $\mathbf{p} - \mathbf{q}$ is perpendicular to S^{c-1} .

In particular, the reflection \mathbf{R} is found by requiring the normal component to be

$$\mathbf{R}_N = -\|\mathbf{R}_N\| \frac{\mathbf{V}_N}{\|\mathbf{V}_N\|} = -\|\mathbf{L}_N\| \frac{\mathbf{V}_N}{\|\mathbf{V}_N\|}$$

This aligns \mathbf{R} with the projection of the view vector onto the normal space (Figure 3). The cosine of the angle between \mathbf{R} and \mathbf{V} is easy to compute.

$$\begin{aligned} \mathbf{V} \cdot \mathbf{R} &= (\mathbf{V}_T + \mathbf{V}_N) \cdot (\mathbf{R}_T + \mathbf{R}_N) \\ &= \mathbf{V}_T \cdot \mathbf{R}_T + \mathbf{V}_N \cdot \mathbf{R}_N \end{aligned}$$

The inner terms of the expansion are zero because the tangent and normal spaces are orthogonal. Substituting for the components \mathbf{R}_T and \mathbf{R}_N yields the specular term for large codimensions:




Figure 4. The light vector \mathbf{L} reflects off a tangent line \mathbf{T} in 3-space, forming a cone of reflections R . For a given view vector \mathbf{V} , the closest unit reflection \mathbf{R} has a component in the normal space \mathbf{N} which is aligned with \mathbf{V} 's component lying in \mathbf{N} .

$$\begin{aligned}
\mathbf{V} \cdot \mathbf{R} &= \mathbf{V}_T \cdot \mathbf{L}_T - \mathbf{V}_N \cdot \|\mathbf{L}_N\| \frac{\mathbf{V}_N}{\|\mathbf{V}_N\|} \\
(2) \quad \mathbf{V} \cdot \mathbf{R} &= \mathbf{V}_T \cdot \mathbf{L}_T - \|\mathbf{V}_N\| \|\mathbf{L}_N\|
\end{aligned}$$

It is convenient to use $\mathbf{V} \cdot R$ to denote the dot product between \mathbf{V} and the nearest unit vector in R .

Even when the codimension of the manifold is 1, equation (2) ignores which side of the manifold is being illuminated, just like in the case of diffuse lighting (equation 1). When \mathbf{V} is more than 90° away from the reflection space R , the dot product $\mathbf{V} \cdot R$ becomes negative. The fact that it can be negative is unrelated to the codimension, and it is reasonable to clamp it to zero. The rationale is that when $\mathbf{V} \cdot R$ is less than zero, \mathbf{V} is so far from the reflection space that it receives no reflected light at all.

The Phong model for calculating the specular intensity is therefore

$$(3) \quad I_{\text{specular}} = k_s I_{\text{source}} (\text{clamp}(\mathbf{V} \cdot R))^{\text{power}}$$

where $\text{clamp}(x)=0$ when $x<0$. The principles for specular illumination are thus (1) light travels in paths of locally minimal length; and (2) the specular reflection is maximized exactly when the view vector nears the reflection space.

This seems like a lot of effort to expend just to end up with the same equation used by Hanson [Hanson93]. But the purpose of this derivation was not to replace the equations. The purpose was to replace the descriptions “ad hoc” and “heuristic” by means of a physically-motivated derivation of the geometric behavior of light, arguing from principles independent of any particular dimension.

3 Compensating For Large Codimensions

When the diffuse model is applied to a k -manifold in n -space, under different values of k and n , a curious phenomenon occurs: the overall brightness increases with the codimension. The torus T^2 is a convenient test object for demonstrating the effect. The surface can be imbedded in 4-space as the cross-product of two circles by the parametrization

$$(x, y, z, w) = (r_1 \cos \theta, r_1 \sin \theta, r_2 \cos \phi, r_2 \sin \phi)$$

where r_1 and r_2 are the “outer” and “inner” radii. One can wrap a curve around T^2 by letting $\phi = A\theta$ for some constant A . The curve or surface can be illuminated in 4-space, or else projected to 3-space and then illuminated there. Illustration 1 shows the result. Notice, especially, how uniformly bright the case $k = 1, n = 4$ is.

In order to understand this phenomenon, first suppose there are light sources uniformly distributed in all directions. How bright is a point \mathbf{p} on the surface or the curve? The answer requires integrating the illumination term over all directions of incoming light. In n -space, these directions cover the unit $(n-1)$ -sphere.

3.1 Surface in 3-space

To integrate the uniform illumination of a point on a surface, let the tangent space \mathbf{T} be the xz -plane and let the light vectors fill a unit sphere. The sphere S^2 has the following parametrization and area element dS^2 .

$$(x, y, z) = (\sin\phi \cos\theta, \sin\phi \sin\theta, \cos\phi)$$

$$dS^2 = |\sin\phi| d\phi d\theta$$

The total area $A(S^2)$ of the sphere is 4π . The area-averaged diffuse illumination $I^{2,3}$ at \mathbf{p} (with $k=2, n=3$) is given by

$$I^{2,3} = \frac{1}{A(S^2)} k_d I_{source} \int_{\mathbf{L} \in S^2} \sin a(\mathbf{L}, \mathbf{L}_T) dS^2$$

The constants k_d and I_{source} will clutter the ensuing calculations; it is convenient to just ignore them (by assuming they are both equal to 1, say). The rest of the computations follow this convention.

Evaluating the integral requires finding an expression for $\sin a$. It is easier to first find $\cos^2 a(\mathbf{L}, \mathbf{L}_T) = \mathbf{L} \cdot \mathbf{L}_T / \|\mathbf{L}_T\|$ by using the dot product. If $\mathbf{L} = (x, y, z)$ then $\mathbf{L}_T = (x, 0, z)$. The sine can be computed from the cosine as follows.

$$\cos^2 a(\mathbf{L}, \mathbf{L}_T) = 1 - \sin^2\theta \sin^2\phi$$

The total illumination for a point on a surface is therefore

$$I^{2,3} = \frac{1}{4\pi} \int_{\phi=0}^{\pi} \int_{\theta=0}^{2\pi} |\sin\theta \sin\phi| |\sin\phi| d\theta d\phi = \frac{1}{2}$$

3.2 Curve in 3-space

Compare the value $I^{2,3}$ to the average illumination of a point on a 1-dimensional curve whose tangent lies in the $(0, 0, 1)$ -direction. The area-averaged illumination $I^{1,3}$ is given by the integral

$$I^{1,3} = \frac{1}{A(S^2)} \int_{\mathbf{L} \in S^2} \sin a(\mathbf{L}, \mathbf{L}_T) dS^2$$

The light's tangent component is $\mathbf{L}_T = (0, 0, z)$, so the sine can be easily calculated from the cosine.

$$\cos a(\mathbf{L}, \mathbf{L}_T) = |\cos \phi|$$

$$\sin a(\mathbf{L}, \mathbf{L}_T) = |\sin \phi|$$

The total illumination for a point on a curve is therefore

$$\begin{aligned} I^{L,3} &= \frac{1}{4\pi} \int_{\phi=0}^{\pi} \int_{\theta=0}^{2\pi} |\sin \phi| |\sin \phi| d\theta d\phi \\ &= \frac{1}{4\pi} 2 \int_{\phi=0}^{\pi/2} \int_{\theta=0}^{2\pi} \sin^2 \phi d\theta d\phi = \frac{\pi}{4} \approx 0.785 \end{aligned}$$

The point is nearly 60% brighter just because the curve has a lower dimension than the surface.

$$\begin{aligned} I^{L,3} &= \frac{1}{4\pi} \int_{\phi=0}^{\pi} \int_{\theta=0}^{2\pi} |\sin \phi| |\sin \phi| d\theta d\phi \\ &= \frac{1}{4\pi} 2 \int_{\phi=0}^{\pi/2} \int_{\theta=0}^{2\pi} \sin^2 \phi d\theta d\phi = \frac{\pi}{4} \approx 0.785 \end{aligned}$$

3.3 Curve in 4-space

If the curve is in 4-space, the point becomes brighter still. The 3-sphere S^3 has the following parametrization and volume element.

$$\begin{aligned} (x, y, z, w) &= (\sin \chi \sin \phi \cos \theta, \sin \chi \sin \phi \sin \theta, \sin \chi \cos \phi, \cos \chi) \\ dS^3 &= |\sin \phi \sin^2 \chi| d\theta d\phi d\chi \end{aligned}$$

The total “surface area” $A(S^3)$ of the 3-sphere is $2\pi^2$. If the tangent is aligned with the $(0, 0, 0, 1)$ direction, the uniformly-lit point \mathbf{p} has an area-averaged intensity which is calculated as follows.

$$\begin{aligned} I^{L,4} &= \frac{1}{A(S^3)} \int_{\mathbf{L} \in S^3} \sin a(\mathbf{L}, \mathbf{L}_T) dS^3 \\ &= \frac{1}{2\pi^2} \int_{\chi=0}^{\pi} \int_{\phi=0}^{\pi} \int_{\theta=0}^{2\pi} |\sin \chi| |\sin \phi \sin^2 \chi| d\theta d\phi d\chi = \frac{8}{3\pi} \approx 0.849 \end{aligned}$$

Similar calculations show that $I^{L,2} = 2/\pi \approx 0.673$ (a curve illuminated in 2-space), and that $I^{2,4} = 2/3 \approx 0.667$ (a surface illuminated in 4-space).

3.4 Exponentiating the Sine

Why does the average intensity increase with the codimension? Consider a k -manifold in $2k+1$ space. For most light vectors \mathbf{L} , the $(k+1)$ -dimensional normal space is closer to \mathbf{L} than the smaller k -dimensional tangent space is. Light vectors that are in, or near, the normal space make a point look bright, so most light vectors reflect brightly when the codimension is large.

It is not enough simply to adjust the diffuse coefficient k_d to compensate for the codimension. Consider what it means for the average illumination to approach the limit of 1: the integrand is bounded above by 1, so it must in fact attain that bound almost everywhere. In almost every direction that light shines, it brightly illuminates almost all of the manifold.

Probably no one is very interested in illuminating a flat object using infinitely many point-light sources distributed uniformly in all directions. The typical situation is complementary to it: there may be a single light source, but the manifold's tangents vary continuously over many (if not all) directions. The visual result is generally the same as the theory predicts: a manifold becomes more uniformly bright when its codimension increases.

A simple way to increase the contrast is to exponentiate using a power $p(k, n)$. This exponent compensates for the surfeit of diffuse reflection. By modifying the diffuse term to be

$$(5) \quad I_{comp} = k_d I_{source} \sin^{p(k, n)} a(\mathbf{L}, \mathbf{L}_T)$$

the brightness is balanced so that a k -manifold in n -space approximates the contrast displayed by a surface in 3-space. The only difficulty is in choosing a suitable value of the exponent $p(k, n)$. It is natural to choose a standard of $p(2, 3) = 1$ since surface-shading in 3-space is the paragon of visual comprehension. For other values of k and n , one proceeds by comparing the averaged integrated intensities $I^{k, n}$ to the averaged integrated intensities $I^{2, 3}$ under the new compensating model of equation (5), finding a value of $p(k, n)$ that makes them equal. The integration is somewhat laborious even for low dimensions, so it is relegated to the appendix. The results are summarized in Table 1, and are applied in Illustration 2.

| | $n = 2$ | $n = 3$ | $n = 4$ |
|---------|---------|--------------|--------------|
| $k = 1$ | $p = 2$ | $p = 4.7635$ | $p = 7.6737$ |
| $k = 2$ | _____ | $p = 1$ | $p = 2$ |

Table 1. Values of the power p used by equation (5) for compensated diffuse illumination of a k -manifold in n -space.

4 Mixing Dimensions

This section describes how the large-codimension model for illumination (equations 1 and 2) can be used to render anisotropic reflectors and furry surfaces. These two examples exhibit a mix of diverse codimensions: 1-dimensional grooves on a 2D surface, or 1-dimensional fibers protruding from a 2D surface. The solution presented here applies to other combinations of codimensions as well.

A manifold may be supplied with one or more vector fields. For example, an isosurface of constant pressure in a fluid flow might possess 1D velocity vectors at every point together with a 2D tangent plane. If vector spaces of different dimensions are associated with a point, one is free to select which space will participate in the lighting calculation. In the case of the pressure-surface, the tangent space reflects like a 2-manifold, whereas the velocity field reflects like a 1-manifold at each sample point.

4.1 Inheritance of Self-shadowing

Section 2.2 noted that a manifold of codimension 1 enjoys the special property of possessing, at most, two sides. The local illumination model can thus simulate the “global” effect of self-shadowing. Consider a surface M with a 2D tangent space \mathbf{T} and a 1D vector field \mathbf{V} in 3-space. A point \mathbf{p} in M is in shadow if its outward normal aims away from the light source. With the light vector directed away from the source, $-\mathbf{N} \cdot \mathbf{L}$ is negative for a self-shadowed point. Assigning a unique normal vector is only possible when \mathbf{T} (of codimension 1) governs the illumination, not \mathbf{V} (of codimension 2). That is unfortunate when one desires to use \mathbf{V} , since self-shadowing enhances the fidelity of a rendered image.

The remedy is to let \mathbf{V} inherit the information (namely $\mathbf{N} \cdot \mathbf{L}$) that informs the model of self-shadowing. To illuminate \mathbf{p} using \mathbf{V} , the reflection terms arising from \mathbf{V} are conditioned by the clamped cosine term arising from the 2-dimensional space \mathbf{T} :

$$(6) \quad I_{conditioned} = (\text{clamp}(-\mathbf{N} \cdot \mathbf{L}) (I_{diffuse} + I_{specular}))$$

The vector space of larger dimension is consulted in order to modify the illumination of a vector space of smaller dimension.

Illustration 3 shows various renderings of a sphere in 3-space endowed with a vector field \mathbf{V} which is tangent to the sphere and aligned in “north-south” directions. Intuitively, this is like a satin ball used as a Christmas ornament. The satin fibers are the 1-dimensional integral curves through \mathbf{V} . The material properties are defined by the coefficients $k_{ambient} = 0.1$, $k_d = 0.5$, $k_s = 1.0$ and an intrinsic color $(r, g, b) = (1.0, 0.25, 0.30)$. In the first image, A, the Kajiya-Hanson model is applied to \mathbf{V} . The rendered curves are just integral curves through the vector field. In the second

image, B , the surface is illuminated using the 2D tangents \mathbf{T} and applying the clamp function. The third image, C , shows the result of illuminating according to \mathbf{V} and interpolating the result over the polygon mesh. The fourth image, D , shows the result of conditioning the solution of image C according to equation (6).

4.2 Attenuation by a Vector Field

Equation (6) shows how the tangent space can be used to simulate global effects in illuminating the 1D vector space over \mathbf{p} . The situation can be reversed as well. In the “satin ball” example, each fiber $\mathbf{V}_{\mathbf{p}}$ was contained in the tangent space $\mathbf{T}_{\mathbf{p}}$ at each point \mathbf{p} . That is, each fiber was constrained to fit the underlying surface. But that need not be the case. Real, physical fibers may protrude outward from a surface, partially shadowing the surface from light. It is possible to simulate this global effect by attenuating the light that reaches $\mathbf{T}_{\mathbf{p}}$. A simple model for attenuation requires the incoming energy to decay exponentially with the distance that it passes through an absorbing medium of density ρ (ρ being between 0 and 1). The medium is the vector field \mathbf{V} . The light generally passes through the medium twice: once on the way in, and again on its reflected path back out. In either case, the distance that it passes through the medium is given by

$$(7) \quad d = h / \sin a$$

where h is the height (perpendicular to $\mathbf{T}_{\mathbf{p}}$) of a fiber at \mathbf{p} and a is the angle between $\mathbf{T}_{\mathbf{p}}$ and the light (entering) or between $\mathbf{T}_{\mathbf{p}}$ and the eye (exiting). The attenuated light therefore has energy I_{atten} given by

$$(8) \quad I_{atten} = I_{source} (1-\rho)^d$$

Illustration 4 shows how conditioning and attenuating the illumination of a mixed-dimensional object can yield convincing results on a torus with radii $r_1 = 1.5$, $r_2 = 0.75$. The fibers have a material property defined by the coefficients $k_{ambient} = 0.1$, $k_d = 0.9$, $k_{spec} = 0.1$, $\rho = 0.02$ and an intrinsic color $(r, g, b) = (1.0, 1.0, 1.0)$ (white). In image A , the vector fields are individually illuminated according to the local model of equation (1). In image B , the vectors are conditioned according to equation (6), using the surface normals of the underlying torus. In image C , the light at the base of each vector is both conditioned and attenuated. Light at the tip is conditioned only. Each vector is shaded as a linearly-interpolated segment. The difference between B and C is especially visible in the lower halves: the fibers blend together in B but are individually visible in C .

5 Running Time

The most complex image (Illustration 4C) contains 409,600 line segments. The fibers are procedurally generated from jittered interpolated samples on the 128×128 mesh of the torus. The image was rendered on a Silicon Graphics Indigo² (75MHz R4400 MIPS processor, 16KB caches,

and 128_{MB} memory) with Extreme graphics, which draws about 150,000 de-aliased Gouraud-shaded vectors per second. The image required 2.2 seconds to compute the illumination on the mesh, 8 seconds to generate the fibers, and 2.5 seconds to draw the fibers. This compares very favorably to image-order (ray-traced) solutions.

6 Conclusions

The diffuse and specular reflection of a k -manifold in n -space can be derived by appealing to four basic principles:

- The re-radiated light's intensity varies with the energy delivered by the incident beam;
- The manifold re-radiates isotropically;
- Light travels in paths of locally minimal length; and
- The specular reflection is maximized exactly when the view vector nears the reflection space.

The resulting equations need to be modified in certain conditions. First, a large codimension generally results in a uniformly-bright object. This effect is ameliorated by exponentiating part of the diffuse term. The exponent can be found via an averaged integration so that the total contrast matches that of a surface in 3-space. Second, a codimension-1 vector space can condition the illumination of an associated vector space of smaller dimension. This permits the smaller space to exhibit self-shadowing. Third, light may be attenuated as it passes through one vector space to reach another. These effects can combine to create effective images of anisotropic and furry surfaces in 3-space. They may also be applied to visualize vector fields over manifolds of higher dimension in large-dimensional spaces.

Appendix

This section calculates the area-averaged compensated illumination for various combinations of k and n . This amounts to solving

$$I^{k,n}_{comp} = \frac{1}{A(S^{n-1})} \int_{\mathbf{L} \in S^2} \sin^p a(\mathbf{L}, \mathbf{L}_T) dS^{n-1} = I^{2,3} = \frac{1}{2}$$

for the exponent $p = p(k, n)$. The definite integral of the exponentiated sine can be looked up in a table: it involves a quotient of gamma functions. Section 3 gives the area element and the volume element for S^2 and S^3 . These quantities have been substituted in the following integrals.

$$I^{1,2}_{comp} = \frac{1}{2\pi} 4 \int_{\theta=0}^{\pi/2} \sin^{p+1} \phi d\theta d\phi = \frac{1}{\sqrt{\pi}} \frac{\Gamma\left(\frac{p+2}{2}\right)}{\Gamma\left(\frac{p+3}{2}\right)}$$

$$I^{1,3}_{comp} = \frac{1}{4\pi} 2 \int_{\phi=0}^{\pi/2} \int_{\theta=0}^{2\pi} \sin^{p+2} \phi d\theta d\phi = \frac{\sqrt{\pi}}{2} \frac{\Gamma\left(\frac{p+3}{2}\right)}{\Gamma\left(\frac{p+4}{2}\right)}$$

$$I^{1,4}_{comp} = \frac{1}{2\pi^2} 4 \int_{\chi=0}^{\pi/2} \int_{\phi=0}^{\pi/2} \int_{\theta=0}^{2\pi} \sin^{p+2} \chi \sin \phi d\theta d\phi d\chi = \frac{1}{2\sqrt{\pi}} \frac{\Gamma\left(\frac{p+3}{2}\right)}{\Gamma\left(\frac{p+4}{2}\right)}$$

$$I^{2,4}_{comp} = \frac{1}{2\pi^2} 4 \int_{\chi=0}^{\pi/2} \int_{\phi=0}^{\pi/2} \int_{\theta=0}^{2\pi} \sin^{p+2} \chi \sin^{p+1} \phi d\theta d\phi d\chi = \frac{\Gamma\left(\frac{p+2}{2}\right)}{\Gamma\left(\frac{p+3}{2}\right)} \frac{\Gamma\left(\frac{p+3}{2}\right)}{\Gamma\left(\frac{p+4}{2}\right)} = \frac{\Gamma\left(\frac{p+2}{2}\right)}{\Gamma\left(\frac{p+4}{2}\right)} = \frac{2}{p+2}$$

To find $p(k, n)$, one merely sets $I^{k,n}_{comp} = 1/2$ and solves for p . The numerical solutions are listed in Table 1.

References

- Anjyo, Ken-ichi, Yoshiaki Usami, and Tsuneya Kurihara. "A Simple Method for Extracting the Natural Beauty of Hair," *SIGGRAPH 92 Proceedings*. Vol. 26, No. 2, pp. 111-120.
- Banks, David C. "Interactive Manipulation and Display of Two-Dimensional Surfaces in Four-Dimensional Space," *1992 Symposium on Interactive 3D Graphics*. ACM Press, pp. 197-207.
- Banks, David C. *Interaction With Surfaces in Four Dimensions Using Computer Graphics*. Doctoral Dissertation. TR93-011, Department of Computer Science, University of North Carolina at Chapel Hill, Chapel Hill, NC 27599. March 1993, pp. 43-48.
- Blinn, James. "Models of Light Reflection for Computer Synthesized Pictures," *SIGGRAPH 77 Proceedings*. Vol. 11, No. 2, pp. 192-198.
- Carey, Scott, Robert Burton, and Douglas Campbell. "Shades of a Higher Dimension," *Computer Graphics World*, October 1987, pp. 93-94.
- Foley, James, Andries van Dam, Steven Feiner, and John Hughes. *Computer Graphics: Principles and Practice* (2nd edition). Addison-Wesley Publishing Company. 1990. Chapter 16.
- Fuchs, Henry, John Poulton, John Eyles, Trey Greer, Jack Goldfeather, David Ellsworth, Steve Molnar, Greg Turk, Brice Tebbs, Laura Israel, "Pixel-Planes 5: A Heterogeneous Multiprocessor Graphics System Using Processor-Enhanced Memories," *SIGGRAPH 89 Proceedings*. Vol 23. No.3, pp. 79-88.
- Hanson, Andrew J. "Interactive Visualization Methods for Four Dimensions," *Visualization 93 Proceedings*. pp. 196-203.
- Hanson, Andrew J. and P. A. Heng, "Visualizing the Fourth Dimension Using Geometry and Light," *Visualization 91 Proceedings*. Vol 21, No. 3, pp 321-328.
- Horn, Berthold and Michael Brooks, ed. *Shape From Shading*. MIT Press, 1989.
- Kajiya, James T. and Timothy L. Kay. "Rendering Fur With Three Dimensional Textures," *SIGGRAPH 89 Proceedings*. Vol. 23, No. 3, pp. 271-280.
- Kajiya, James T. "Anisotropic Reflection Models," *SIGGRAPH 85 Proceedings*. Vol. 19, No. 3, pp. 15-21.
- LeBlanc, A., R. Turner, and D. Thalmann, "Rendering Hair Using Pixel Blending and Shadow Buffers," *Journal of Visualization and Computer Animation*, Vol. 2, No. 3, August 1991.
- Miller, Gavin. "From Wire-Frames to Furry Animals," *Proceedings of Graphics Interface 88*. pp. 138-145.
- Poulin, Pierre and Alain Fournier. "A Model for Anisotropic Refelection," *SIGGRAPH 90 Proceedings*. Vol 24, No. 4, pp. 272-282.
- Rogers, David F. *Procedural Elements for Computer Graphics*. McGraw-Hill Book Company, 1985. Chapter 5.

Steiner, K. Victor and Robert Burton. "Hidden Volumes: The 4th Dimension," *Computer Graphics World*, February 1987, pp. 71-74.

Ward, Gregory J. "Measuring and Modeling Anisotropic Reflection," *SIGGRAPH 92 Proceedings*. Vol. 26, No. 2, pp. 265-272.

Watanabe, Yasuhiko and Yasuhito Suenaga. "A Trigonal Prism-Based Method for Hair Image Generation," *IEEE Computer Graphics and Applications*. January 1992, Vol. 12, No. 1, pp. 47-53.

Westin, Stephen H., James Arvo, and Kenneth Torrance, "Predicting Reflectance Functions from Complex Surfaces." *SIGGRAPH 92 Proceedings*. Vol 26, No. 2, pp. 255-264.